



LIBRARY
OF THE
MASSACHUSETTS INSTITUTE
OF TECHNOLOGY

TOWARD THE REALIZATION OF INTELLIGENT
MANAGEMENT INFORMATION SYSTEMS*

by

Donald C. Carroll** and Zenon S. Zannetos***

218-66

October 10, 1966

*Prepared for presentation at the Third Congress on the Information Systems Sciences, Sponsored by the Air Force Electronic Systems Division and The MITRE Corporation, Buck Hill Falls, Pennsylvania, November 21-23, 1966.

**Associate Professor of Management, M.I.T., and MITRE Corporation.

***Professor of Management, M.I.T.

TOWARD THE REALIZATION OF INTELLIGENT
MANAGEMENT INFORMATION SYSTEMS

Donald C. Carroll and Zenon S. Zannetos
Sloan School of Management
Massachusetts Institute of Technology

I. INTRODUCTION

Statement of Purpose

Our purpose is to propose some new doctrine for management information system design, to state some explicit goals to be sought, and, in so doing, to offer some new perspectives for designers. We will also briefly review the current state of the art in light of this doctrine and propose some steps towards realization of the proposed goals.

In reviewing existing statements on what a management information system should provide, we have noted a singular lack of operationally viable goals. "To provide a basis for better decision making" simply does not provide a basis for choice for the system designer. It is to help fill this void that we are motivated.

On the Nature of Intelligence

In a nutshell, what we are proposing is an information system for increasing management's understanding of its environment and its rationality

in dealing with it. As we have considered what underlies such a system, we have been struck with the similarity between our proposal and the properties attributed by students of cognitive process psychology and heuristic programming to "intelligence".¹ That is, what we are suggesting are information systems which will allow organizations to exhibit more intelligent behavior.

To provide a rigorous definition of "intelligent behavior" is extraordinarily difficult, as witness the long-standing controversy over criteria for machine intelligence.² Most scholars now resort to "Turing's test"³ to settle the question of machine intelligence; others define properties of intelligence, capabilities which are apparently necessary to behavior which is generally accepted as intelligent. For example, in Minsky's "Steps Towards Artificial Intelligence", there are identified several basic processes such as non-random or non-exhaustive search, pattern recognition, "learning" (in a limited sense), "planning" (abstraction is the central idea he ascribes), and induction, as being associated with intelligent behavior [28]. We shall define the elements of our doctrine in terms of similar properties.

¹Particularly Minsky [28], and Simon and Newell [30].

²For a review see [2].

³The basic idea of the test is remotely to connect the examiner (by teletype, for example) to a computer and to a person equipped also with a teletype. If the examiner cannot distinguish which responder is the man and which is the machine, the machine is "intelligent" (as intelligent as the competing man in the context of the particular discussion one assumes). See [39].

But with perhaps more courage than wisdom, we do wish to offer a working definition of intelligence, if only to clarify our use of the term as applied to management information systems. Intelligent behavior is "doing what you do for the right reasons", roughly speaking. This has two facets: first, understanding of cause and effect, and second, rationality in employing that understanding. Note that we haven't said "doing what is right". Our definition admits of intelligent behavior in the face of uncertainty; it requires neither omniscience nor clairvoyance, and distinguishes between the decision-making process and the outcome of such.

Let us be more specific about this concept of "understanding". The better one understands a phenomenon, the more accurately he can predict its behavior. And, for purposes of subsequent discussion we shall find it convenient to associate understanding with a predictive model of a phenomenon. That is to say, by increasing managerial understanding, we shall mean improving the manager's model of the real world -- his model whether it be explicit or implicit, being his basis for prediction.⁴

The rationality aspect of intelligence is relatively clear; given a set of relationships and facts, one can be more or less adroit in exploiting this knowledge for his own ends. There are at least two types of situations in which logic is at least as important as knowledge. One is in decision making under uncertainty. Here statistical decision theory is now offering a suitable

⁴Pounds [31] has made a strong case for the universality of models in managerial behavior, in "problem finding" as well as in problem solving.

rationale (although in problems of large size, the computational side of the logic becomes important). The second situation is in system management. The detailed interrelationships among system components may be well understood but overall system behavior is not, either because of dynamic effects or simply because of the combinatorial nature of the full system. Forrester [19] now has offered an approach in the former case, and simulation or heuristic programming some promise for the latter.

We should also stress that intelligence is both a relative and dynamic quality of persons (and organizations). We cannot, except arbitrarily, divide behavior into intelligent and unintelligent, but we can talk about "levels" of intelligence. We are willing to admit that von Neumann was smarter than we, individually and collectively (we have some suspicions about some others as well). Hence, we shall be concerned with people or organizations who are more or less intelligent. In stating that intelligence is dynamic, we simply mean that one can grow (or even regress) in the qualities associated with intelligent behavior. In other words, one can learn.

In the light of the above, we can now be somewhat more precise in our goal: we want information systems designed to provide for more intelligent behavior over time.

We request license to attribute intelligence to a group, or an organization, and in so doing, we recognize that there will always be great variation in the intelligence exhibited in different aspects of the organization's behavior -- with a group considered collectively there is a guaranteed multiple split personality, so to speak. However, the legal system prosecutes the corporation

as an individual, and poets and historians have frequently treated armies such as Alexander's or Napoleon's as having but one (or, at most one) intellect. And finally, our treatment of multiple parallel processors as a single parallel processor (i.e., an organization as one person) is probably less heroic than the treatment of one parallel processor as a sequential processor as perpetrated by researchers in machine intelligence.⁵

On the Nature of Management

Our arguments are based on some basic assumptions about the processes of management. Our main assumption is this: a management organization, based on its collective understanding of causal relationships among resource inputs, processes, external environmental factors and outputs, manipulates those things under its control to achieve results in accord with its objectives.

In our view, the focal points for those aspects of management which can be considered cognitive are planning and control.⁶ Planning, as we construe it, subsumes all consideration of the future course of events, whether the time scale is short or long, whether the process is formal or informal, and whether its inputs and outputs are physical entities or mental processes. Planning thus includes searching for future alternative courses of action, selection of goals, specification of procedures to be followed or resources to be acquired and utilized for the achievement of the chosen goals. We distinguish between operational planning, in which the emphasis is on what the organization should be

⁵A problem noted by Selfridge and Neisser [36] as well as by Newell and Simon [29].

⁶Our classifications, planning and control, are quite inclusive. They correspond roughly with Anthony's [1].

doing in the relatively short run as constrained by the dominant characteristics of its current structure, and strategic planning, wherein the emphasis is on possible changes in the dominant characteristics of the structure (physical or organizational) or in major goals. The nature of the planning process entailed is likely to be quite different in the two types, to wit: operational planning is typically a continuing, systematic process whereas strategic planning is more often ad hoc and unstructured; operational planning is intimately linked with control, providing milestones or other goals and drawing from the control system current status, whereas strategic planning is likely to consist of one-shot ("terminal") decisions only loosely linked to the formal control system.

Regardless of emphasis, planning always involves a model; the model is explicit in many cases, but is certainly implicit in any activity that projects the future.⁷

The object of control is to obtain desired behavior (often as set forth in a plan). Control operates after the fact of execution.⁸ There are certain basic processes which go into control. The first is measurement of the status or performance of the controlled entity. But the measurement has no meaning until it is juxtaposed with the standard or desired measurement (in general, the "set point"), consequently comparison is a basic process. The third process

⁷For a thorough discussion see [14].

⁸We are restricting ourselves to so-called "closed loop" control systems as opposed to "open loop" or "calibrated" control systems. We have not seen any of the latter employed in management.

is direction in which the controlled activity receives signals to alter its behavior to obtain closer conformance with that desired.⁹

Measurement, comparison, and direction are common to all feedback control, but quite often there are additional subprocesses present. For example, in system control, one measurement may relate to several different activities or entities (this is part of the notion of "integrated data processing"). The fact of completion of a task in a manufacturing shop may be reflected in individual worker productivity, the foreman's direct labor expenditure, production progress control, and the like. So within the control cycle there is required a subprocess of classification, that is, association of the measurements with the appropriate "responsible" entities.

Also, implicit in the comparison-direction portion of the loop is determination of the cause of deviation of actual measurement from desired. Diagnosis is the term we will use. Now, diagnosis is a trivial process in some instances, especially in automatic process control. In a household heat control system, the furnace is always assumed to be the culprit when actual temperature departs from desired, at least for purposes of direction. On the other hand, in a process of any complexity, particularly when probabilistic elements such as human beings play a part, diagnosis can be an extremely complex process. A time overrun on an activity in a PERT network seldom has a simple cause, for example. In manufacturing cost and schedule control systems, allocation of blame between performer and standards estimator has

⁹There are situations in which deviations of one sign only are "bad", deviations of opposite sign therefore receive no correcting signal; if anything, they receive "reinforcement".

always been a source of argument, to say the least. Diagnosis may occur before direction, in which case causality is a consideration in formulating the direction, or it may occur as a result of the direction, in which case it is performed by the controlled entity.¹⁰

Control systems differ as to the specificity of the desired behavior. In simple cases, the purpose of control is strictly regulative, keeping performance within reasonable limits. But in other cases, again especially when people are involved, the control system assumes an educative role. That is, the control signals encourage desirable aspects of activity by "rewarding" them in some way (or by "punishing" undesirable aspects), thus leading in theory to process improvement, "learning" again.¹¹ All incentive systems, whether applied to top executive or to workers are educative in purpose.

Control, just as planning, is always based on a model. The model may be a simple expression of formal cause and effect relationships (e.g., furnace yields heat, bad supervision yields unfavorable labor variance), or it may be highly informal, implicit in post mortem analyses effected upon major deviations, or it may be an explicit mathematical model.¹² We shall

¹⁰It may be useful to conceive of the former as autocratic diagnosis and the latter as democratic. We will not venture a value judgment on the relative efficacy of the two possibilities.

¹¹We should stress the generality of our concept of reward and punishment. Simply passing on information suggestive of good behavior (as in democratic control) is a reward in this sense.

¹²These are causal models. Of course, the "set point" or other statement of desired behavior is a model, as well.

later attempt to classify control systems on the basis of the sophistication of the control model.

Hierarchy in Planning and Control

All management organizations are hierarchical to a degree,¹³ and this implies hierarchy in the processes of planning and control. Even at one organizational level planning is "senior" to control in the sense that it provides the "set point" for the control process. Another potential interrelation between planning (particularly operational planning) and control occurs in the "diagnostic" process. As the control model is revised, presumably the planning model should reflect the new insights. We shall advocate this.

There are several levels of planning and control in a management organization. We expect to see at the lowest level of the system detailed plans (often in the form of schedules) driving the control of the basic productive physical processes. At higher levels, we expect to see efforts to coordinate (via a plan) the control of multiple interdependent activities. Lower level managers control operating processes, but higher level managers control lower level managers, it has been noted. Anthony, in fact, draws a sharp distinction between "operational control" and "management control" (in the latter he is referring to the manager as the controlled entity as well as the controller) [1]. We will be particularly concerned with a related aspect, that of controlling the planning process. Consequently, we will draw a sharp

¹³For a thorough discussion, see [43].

distinction between planning process control and operating process control.¹⁴

On the Virtues of Intelligent Management

While campaigning for intelligent management might appear on the face of it to be about as controversial as advocating motherhood and good works, some consideration of the economic justification of our particular forms of intelligence ought to be undertaken. We seek better understanding, models which more faithfully represent the real world, and we seek rationality, better logic applied to this understanding. A more valid model means more predictable execution of plans as previously noted. Better logic in developing plans means better alternative courses of action selected, other things being equal. In regulative control, a better model enables reduced variation in performance -- output to closer tolerances, for example, and in educative control, a better model enables what the learning theorists call more accurate "discrimination" of desirable behavior and hence, more rapid improvement.

The nature of discrimination can be illustrated by the story of B. F. Skinner's "superstitious pigeons" [23, p. 88]. Several pigeons were placed in separate boxes. A feeding mechanism delivered food to each pigeon every fifteen seconds regardless of what the pigeon was doing. After operating in this way for some time, the experimenter observed that one bird was sitting very still, another bowing, another turning around in tight circles, another

¹⁴We are well aware of even higher levels such as planning process control planning and planning process control control, and so forth. We will later assert that with regard to the basic information system, that which is sufficient for planning process control is sufficient for higher levels as well.

hopping on one foot, and so on. Each bird repeated its own ritual between feedings. Analogous (presumably civilized) human reaction to random rewards or punishment are frequently found in competitive athletics and, we suspect, in management.

We will offer subsidiary arguments for formality in these planning and control processes. Among our several reasons is this. While various cognitive agents, namely, people, may come and go, for the organization to increase its intelligence over time, it must make some provision for recording the accumulated planning knowledge and thus guard against loss of memory together with the people. A formal model can be stored in "memory" and hence provide continuity in intellectual growth.

Intelligence Revisited

We have noted the presence of certain processes in managerial planning and control, namely: measurement, classification-association, search (in planning), learning (in educative control), and diagnosis. The latter, as we shall discuss, can involve the most complex forms of pattern recognition, abstraction and inductive inference. We have thus found relevant to our doctrine the central processes of intelligent behavior.

We will choose as the focal point of our discussion, the feedback-comparison process of control because the opportunity to improve the environmental model starts when behavior and the model diverge. But recall that we will be discussing both operating process control and planning process control. In the latter, we will be considering an essentially

inductive process (diagnosis-model improvement) applied in analysis of an essentially deductive process (planning).

What Lies Ahead

In Part II of this paper we attempt to establish desirable operational characteristics of intelligent organizational behavior and to translate these into requirements for the information system. We review the current state of the art with an eye towards identifying instances where these requirements are being met (in part) in Part III. In Part IV, we propose some solutions and research to be undertaken to meet these requirements. Part V is devoted to a brief summary and broad conclusions.

II. SPECIFICATIONS FOR INTELLIGENT INFORMATION SYSTEMS

Introduction

We are focusing upon operating process control and planning process control as the central cognitive processes of a management organization. In operating process control, we wish mainly to obtain specified behavior of the activities which comprise operations. In planning process control, we assume primarily an educative goal, that is to say, we wish explicitly to improve the planning process. We will first treat the problem of regulation and improvement in operations and identify therein the potential for structures of different "levels of intelligence", and we will attempt to specify the information system requirements implied for the highest defined level of intelligence. We will then turn to the more complex problem of planning process control. The difficulties here stem from less tangible process goals, less formal processes, and unclear boundaries on the problems being attacked. The control process is the same in both cases, but the process being controlled is sufficiently different that only a highly intelligent control system will suffice to assure improvement, it will be argued.

Operating Process Control

There is always a process model which underlies control. It is on the basis of this model that the magnitude and sign (in general, the nature) of the control direction is determined, and, in more complicated systems, that the particular agency to receive the signal is determined.

The model can be naive or sophisticated. This is partially a question of the complexity entailed in the model, but in our view, more fundamentally related to the depth to which cause and effect relationships are captured. Applying a polar classification to a continuum, we identify the extremes as symptomatic and causal control. That there is a continuum of causality should be clear to anyone who has attempted to respond in good faith to a three-year-old's infinite series of "whys". An example of clearly symptomatic control is a wage incentive system used for educative productivity control. Output and reward are directly linked, and little formal attention is paid to causes of output except when major dislocations such as machine breakdowns or material shortages disrupt the process. The assumption is made that high output results from energetic or skilled attention to duty by the worker. If this assumption is largely correct, the system works. But if output is affected by a substantial number of causes other than the worker's activity, the system can be acrimonious in its administration and ineffective in its application.¹⁵ A more causally oriented control system applied to the same problem would attempt to correct output variations for "degree of difficulty" so to speak, by removing the effects of differences among tasks (i.e., more precise standards), differences among materials or material suppliers, differences among machines and the like. It would, in this case, isolate as nearly as possible that portion of output variation truly attributable to the worker. In the extreme case, it would attempt to

¹⁵The battle among workers for "make out", i.e., tasks for which causes other than the workers' effort make good performance easy, is behavior symptomatic of this type of problem.

classify elements of the worker's behavior as being causally related to output and to reward appropriately, which is to say, it would assist in discrimination.

The close connection between causality and understanding implies that causal control provides a higher level of intelligence than symptomatic.

Another dimension of classification related to the question of intelligence is the adaptivity of the model. At the lowest level in this case is reflexive control, based on a fixed model with fixed or externally supplied parameter values. The name derives from the parallel to a reflex in human physical behavior; the effect is that of a "stored response" to stimuli. Reflexive control is employed in very simple situations such as the home-heating thermostat example cited above as well as in highly complicated system inventory control based on massive mathematical models. The salient point is that the system behavior is not easily modified and is certainly not self-adapting to a changing environment.

Moving up the scale, the next step is parametrically adaptive control. The model is fixed as to the constituency of variables and parameters and the relationships among them, but the values of the parameters are changed as a function of experience or as exogenously supplied data vary.¹⁶ This type of system is often seen now in chemical process control, in which the model relates yield to a variety of input and process factors. The weights

¹⁶Such systems can be thoroughly sophisticated. See Bellman [3].

placed upon the factors are varied as experience accumulates, providing lagged response to new environmental information. Another aspect of parametric adaptivity is seen in inventory control based on "adaptive smoothing".¹⁷ In this case, the smoothing parameters are adjusted as data are processed in an attempt continuously to obtain a minimum variance forecast.

At a still higher level, we can hypothesize inductive control, in which the entire model is subject to change, both in structure and constituency, as well as parameter value. What we have in mind here is continuous re-evaluation of the model, diagnosis being the central process. Inductive control, in attempting to establish causality in greater depth, involves formulation of new hypotheses and tests thereof. Indeed, since it is advancing "hypotheses of causes" it parallels closely the Bayesian "prior to posterior" process with highly complex multivariable models.¹⁸

We would ask an additional step in inductive control. Since the control model is being adapted, it would seem essential to adapt the related planning model as well. In fact, we will denote inductive systems, which provide for direct updating of planning models as part of the same process, as prognostic as well as diagnostic in purpose.

¹⁷Brown [6] has a complete description and discussion.

¹⁸The "hypothesis of causes" was Bayes' own name for his theorem. Basic references in Bayesian statistical decision theory are Schlaifer [35] for the layman and Raiffa and Schlaifer [33] for the expert.

We know of no examples of formal inductive control systems in operation. To clarify the ideas, however, consider the following situation.¹⁹ There was a statistical analysis of yield performed on a mechanical process. The purpose was to relate process yield to various (controllable) input and process factors and ultimately to increase the yield. In due course, a rather elegant predictive model was obtained, but the remaining unexplained process variation was still substantial. Additional variables were studied to little effect until finally the time of day during which the process was being operated was tested. This variable showed almost dominating significance. And, further investigation showed that the third shift superintendent was paying essentially no attention to process yields with predictable effects. While this situation may illustrate either missing the forest due to overzealous tree examination or serendipity (depending on what the responsibilities of the analysts are assumed to have been), it is a clear example of economically significant increased understanding resulting from an inductive control process. We shall later cite additional cases where the purpose of inductive control is implied in informal systems.

Levels of Intelligence in Operating Process Control

As must be clear, we would classify reflexive-symptomatic control at the bottom of our scale and (prognostic) inductive-causal at the top, since the latter provides the potential for achievement of the highest level of intelligence in an organization -- for learning, in the general sense.

¹⁹One of the authors played a role in the situation which is better left undefined.

But let it be clearly understood that we are not advocating wholesale redesign of all operating process control systems to achieve this sort of intelligence. In fact, to the extent that the environment is stable and very well understood, a reflexive control structure may be wholly adequate. After all there are many cases where response to the symptoms also cures the disease. To the extent that constituency and general functional relationships are well understood, parametrically adaptive control may be just right. On the other hand, to the extent that the environment is not perfectly understood or is changing, then there exists an argument for inductive control.²⁰ As we view the world, the latter category appears to include the majority of systems-control problems and the majority of activities subject to rapid technological, political, or market change.

We therefore admit to exceptions but advocate intelligent control as a rule.

Information System Requirements for Intelligent Operating Process Control

An initial step in establishing causality is to establish association between the basic variable or variables of interest and other factors capable of being measured and either corrected for or controlled themselves. Thus, the first step in uncovering the causes of lung cancer has been to establish the association of the incidence of that disease with cigarette

²⁰There is always a question, too, of the economically justifiable depth of diagnosis in inductive control. Since causality is not necessary for predictability (it is sufficient), it may be optimal to cease searching at some symptomatic level.

smoking. Association is necessary for causality but not sufficient to prove it; it is a first step.²¹ Since diagnosis occurs after the fact, as with other post mortem (or post victorian) activities, it begs for recreation, in a flexible way, of the situation when the unexpected occurred. This implies a requirement for a variety of associations among factors, temporal as well as functional, for adequate feeding of the diagnostic process. Furthermore, a difficulty in establishing association is confounding, the inability to separate the effects of two or more variables due to overly gross or aggregated measurement. Hence, we require in our supporting information system the facility for functional and temporal association with precision and in detail.

The problem of deciding just which measurements should be maintained is difficult. Potentially relevant data, not just known relevant data, are needed, if the model itself is to be modified. This fact may explain the popularity of parametrically adaptive control models; with them at least the data base is well-defined.

Let us attempt to be more specific about the idea of association. What is needed is a way of finding out the values of a large number of variables which were current at some point in time. Functional association requires linkages among the factors and the basic process measurement. In the yield analysis described above, for instance, all of the measurements

²¹With rare exception, causality cannot be established statistically; proof of sufficiency often requires systematic elimination of all other possible causes, or controlled experimentation.

of input and process characteristics had to be linked to the yield on a particular batch. The temporal association capability allows for lagged effects, and for dynamic analysis of phenomena in general.

A detailed associative data base is the raw material for inductive control, but additional capabilities are required for the diagnostic elements. Some aspects that are known to be present (but which are not well-understood) are pattern recognition and pattern generation. The former includes the ability to perceive relevant associations and to match a given pattern to observed behavior. Often this involves "normalizing" the data, putting them in a proper format or otherwise transforming them to conform to the pattern or patterns being tested. For example, simply arraying data in time series form normalizes them for certain dynamic pattern matching; in other cases, a graphing of frequency spectra might be required.

Pattern generation is even less well-understood, but it clearly involves abstraction and quantitative hypothesis formulation, which is to say, model building. The question begged is what is the source of the model. Much opinion suggests that there exists frameworks, general theories, or taxonomies -- broad categorizations of phenomena -- which suggest detailed models for testing. Freudian psychology, Marshallian economic theory, or more recently, Forrester's "industrial dynamics" [19], are examples of formal frameworks. In general, however, the totality of

one's experience, observation, and education serves as the framework for a human. Pounds [31] suggests that the process of diagnosis begins (a "problem is found") when behavior departs from that suggested by one of these frameworks.

An example may serve to clarify what we mean by pattern recognition and generation. Forrester [19] has cited some instances of self-induced oscillatory behavior in business, one of which was evidenced by inexplicable seasonal demand for a consumer product. By drawing on his general framework he was able to construct a model which (qualitatively) matched that of the (normalized, measured) behavior of the firm. From his model, he was able to deduce that the source of the seasonal peaks and valleys were the firm's traditional promotional patterns and its customers' anticipation of this pattern.

Pattern recognition, abstraction and hypotheses remain shrouded in mystery as to their precise mechanisms; they are apparently tied up with the very most arcane human capabilities which are often collectively labeled "creativity".²² The mystery notwithstanding, we require these faculties as operative elements in inductive control. Also bear in mind that they must be employed in the worst of all possible inferential worlds evidencing as it does probabilistic and dynamically non-stationary behavior and imperfect measuring devices.

²² Minsky [28] and Newell and Simon [30] have much to say on this. The fact is that, at this point in time, people can do these things very well and machines not well at all. See also Licklider [26].

Planning Process Control

When the planning process is brought under surveillance, all of the previously cited aspects of control apply, but there are some new problems to face, as well. Some of these can be stated as follows:

1. Planning is always based on a model, so in control of planning there is a metamodeling problem. We require a model of the (planner's) modeling process and a model of the employment of the planning model.

2. Planning, especially strategic planning, is often based on information about matters external to the firm or organization. For example, predictions of competitors' behavior or general economic conditions are often basic to commercial planning. The enemy's order of battle (in general, intelligence²³) occupies a similar position in military planning. Hence, planning process control requires a data base that is not necessarily a convenient by-product of operating process control or otherwise at the disposal of its users.

3. In order to establish control, there must be a process goal or standard, in this case an objective purpose for planning. Yet it is not always abundantly clear just what one is attempting to achieve by planning.

²³G-2 not IQ.

4. Planning is frequently intuitive and subjective both as to process and to data. Yet in order to exert control, the subjective estimates and value judgment require quantification and their functional relationship with the planning goal requires establishment.

5. Planning, in many cases, looks far into the future. It would be desirable to conduct post mortem analysis for process improvement, yet if the planning process controller waits for the future to reveal itself completely, the control cycle time will be too long.

6. Planning is typically a group rather than an individual process. We understand little enough about individual behavior but even less about group behavior. Operating also is often a group process, but planning is "groupthink" rather than "groupdo".

7. Correlatively, planning is a task still (if temporarily) performed largely by people. Attempts to observe or to experiment with people often leads to a well-known phenomenon, known as the "Hawthorne Effect", in which the subjects respond to the fact or conditions of the experiment rather than to the

environment being studied.²⁴ What is even worse, often the environment surrounding the experiment itself changes by the mere fact of being observed.

This list is long enough to be discouraging to the most ardent of idealists; but the alternative of uncontrolled planning should be sufficiently dismaying to make the effort worthwhile. And this list suggests a process sufficiently poorly understood to require inductive-causal control.

Information System Requirements for Planning Process Control

We clearly require an upgraded data base for planning process control. Its scope requires expansion to include both external data (including forecasts of external variables) and subjective data. By the latter, we mean that the planning assumptions, subjective estimates, and value judgments should be formally recorded. And, of course, we require the same associative facility with these data as we did for operating process control.

What this amounts to is a plea for formal models in planning which we add to those previously voiced by others, notably Emery [14]. The added motivation is the potential here for planning process improvement. To this we add the requirement of formal goals for the planning process.

²⁴The name stems from some working condition experiments conducted at the Hawthorne Works of Western Electric in the thirties. A group of women workers was submitted to varying lighting, heating and other factors. Regardless of conditions their productivity rose. The experimenters finally concluded that the women were responding to the attention that accompanied the experimentation. Described in [34].

For control, it is not sufficient merely to evaluate the product -- the plan -- we require access to the process as it operates. This means somehow capturing the "stream of consciousness" of the planner to obtain his "trace", i.e., the logic used to formulate his plan. And, to tighten the loop in long range planning we need a method for analyzing incomplete returns, to infer on the basis of partial data. And, we require as before a powerful diagnostic facility. Finally, some provision, such as clandestine, unexpected, or constant surveillance must be made to avoid the Hawthorne Effect.

Higher Order Processes

We have previously noted the possible existence of higher order processes such as planning process control control. We submit that the information system requirements for these processes are generically no different from those set forth above.

III. COMMENTARY ON THE CURRENT STATE OF THE ART

Introduction

We will attempt to review what we perceive to be examples of elements of intelligent information systems which are generally or specifically in operation. We face a problem in so doing in that we suspect that organizations which have achieved higher levels of intelligence are probably intelligent enough not to publicize the fact, so the state may not be so primitive as we represent it.

The State of Operating Process Control Systems

Control systems for detailed productive processes have been growing in sophistication ever since computers became generally available. In continuous process control, for example, very elaborate formal-model based systems for chemical processing are now common. These vary in their complexity, but most commercially available systems are capable of multi-variable control at multiple levels (i.e., they adjust the process to conform with "set points" on several variables and also compute the proper value for the set point based on external inputs). Parametrically adaptive systems of at least modest scope are operative as well. However, since these systems are fixed as to model structure, such inductive inference towards model improvement as takes place must be performed externally to the system. These systems are of interest as models for man-machine system control, but while sophisticated in model structure, they offer no guide for model improvement -- they do not evolve.

Another area of interesting development is that of detailed job shop production scheduling and control as practiced in Hughes Aircraft [37] and Westinghouse Electric [38], among other firms. The general structure of these systems is this. A simulation model is fed with inputs of the current order backlog (with routings, processing time estimates, and due dates), the shop configuration (machines, and men), and some decision rules for dispatching the jobs. The model is run and rerun, simulating the future course of events, allowing for adjustments to backlog (i.e., subcontracting), or capacity (overtime, added shifts) and the decision rules. When a "satisfactory" simulation is obtained, the simulated start time of each job on each machine is used as the scheduled start time for the job in the shop. This schedule provides the set points for production control.

These "finite capacity" schedulers (so-called because they explicitly consider the availability of the work station before simulating the assignment of a task) are considerably more complex than the standard "infinite capacity" scheduling systems in which a scheduled date for each task is obtained by dating back from the job due date using "standard lead times" (which allow for direct production time, waiting, transit, setup and the like) for each operation on the routing.²⁵

²⁵Emery [13] provides a discussion in depth of various alternative scheduling systems including these two.

One problem with the infinite capacity schedules is that they are not feasible, even in theory. The schedule dates provide crude targets for progress, but because work station capacity is not directly considered, a deviation from schedule may only signify that the schedule was impossible at the outset. The trouble is that the deviations resulting from model inadequacy are confounded with true process deviations. This is less a problem with a finite scheduler. A deviation in this case generally indicates that something unexpected has occurred such as low productivity, a bad processing time estimate, a material shortage, or failure to follow the scheduled sequence. While causality is not pinpointed, a point of departure has been established. (Even in finite capacity schedules, minor deviations tend to compound after a time and schedule infeasibility again rears its head. Potentially, this schedule "decay" can be cured, and discrimination of causes materially improved in on-line, real-time control systems. This possibility is discussed in the next section.)

The more detailed model (derived from the more detailed data base) used in finite scheduling, and the built-in time-based association of resources (work stations) with activities provides the increased control power. In effect, the better model eliminates "noise" from the information system; a deviation signal signifies something. In comparison, the naive infinite capacity schedule tends to "cry wolf", leading to ineffective remedial action. Also noteworthy is the use of a common model for planning and control. To the degree that experience improves the model (such as by

improving processing time estimates or making more precise allowance for setup time), the improvement applies immediately in scheduling. The systems therefore have some prognostic power.

In conventional accounting control systems the state of the art is dismally primitive. Budgets and other standards are frequently almost arbitrarily arrived at, only major deviations may have any significance and they could easily result from factors totally beyond the aegis of the controlled entity. There is no systematic way of filtering noise from the information system and no aids are provided for causal diagnosis, or even determining significance. This is not to say that managers do not attempt to determine causes of budget overruns, for example; it is to say that such diagnosis occurs separately from the control system (and in some case, in spite of the control system). At best, in the absence of managerial brilliance, the conventional wisdom in accounting control amounts to symptomatic-reflexive control, the bottom rung of our intelligence ladder. It is small wonder that managerial behavior approximating that of the "superstitious pigeon" is not uncommon.

There are, however, some candles being lit in this area of stygian darkness. The general idea of the "flexible budget", of separating deviations from plan due to volume ("volume variance") from deviations due to, say, labor performance and use of material resources independent of volume, represents an attempt to separate gross uncontrollable (by the

production manager in this case) effects from those which can properly be laid at his door. But the accountants, having possibly exhausted their creativity, seem less than eager to press on in this direction.

More encouraging are the recent attempts in the Bell System [22] and elsewhere [4, 48] to establish performance standards on the basis of more precise statistical models. For example, suppose there is a multiple plant company with each plant producing comparable products. An electrical utility will serve. There will be variation among plants in measures of performance, say average delivered cost per kilowatt hour. Some of this variation is attributable to plant management, but a great deal to factors outside the control of management, such as fuel costs, age of generating equipment, population density, climate, longitude of the area served, industrial concentration and classification and so on. Merely to compare plants on the basis of the raw performance measure is patently unfair to the plant manager who has drawn unfavorable circumstances. Mean performance as a function of all of the uncontrollable factors can be predicted for a particular plant on the basis of statistical (i.e., multiple regression²⁶) models. This revised performance figure represents a standard calibrated for "degree of difficulty", so to speak. Performance deviations from this calibrated figure represent true "managed" performance plus a much smaller "unexplained" component, and certainly provide a more accurate basis for learning, reward or castigation.

²⁶Broad coverage of the technique can be found in Ezekial and Fox [16] and Graybill [21].

These control systems, in their separation of uncontrollable causes from controllable, represent a major step towards causally oriented operating process control. While the particular techniques used in the cited studies may be limited in their applicability to relatively homogenous product oriented industries, the philosophy underlying the use appears impeccable to us. An interesting possibility for extension would be to associate the residual variance with identifiable management-controlled variables, e.g., work force composition, salary levels, some quantified aspects of operating strategy. Associations of this sort would lead naturally to diagnosis -- inductive control.

We have observed efforts towards diagnosis in system management, particularly in the PERT-based planning and control system employed by NASA in the APOLLO program. First, there is a formal model used for planning and control. Second, the evidently widespread doctrine of "visibility" is employed in project time and cost (and to a lesser degree, technical performance) control. This calls for focusing attention on responsible parties in cases of unfavorable deviation from plan. From discussions with both the system managers and contractors it seems clear that what goes on in the "control rooms" during the post mortem project reviews is causally oriented diagnosis to a substantial degree.²⁷ The significant point is that the whole information system appears to be oriented towards this process. We feel that "visibility" insofar as it encourages diagnosis, is a useful system design concept.

²⁷The first item on the agenda is naturally an inquiry into what can be done to correct the deviation as it exists, only then comes the "Why?".

In addition, several attempts have been made to "calibrate" the bias of the planning estimates of the contractors and thence to correct the plan for this bias as it becomes known, a clearly prognostic exercise.

The State of Planning Process Control

In general, it appears to us that organizations have recognized the need for planning process control for years, but little or no formal surveillance has been instituted. For example, there is usually some control exercised over the process of budget preparation in government and industry, in the form of critiques of assumptions and also end-of-fiscal-year post mortems. One of the clearest examples of this is the Westinghouse Electric "Profit Planning" system described by Evans [15]. High level (product department) plans are examined, reviewed, and critiqued on the basis of their assumptions and substance before the execution begins, and periodically during the year, the execution is reviewed. Care is taken to separate effects due to poor planning or poor forecasting from poor performance both by dialogue and in the structure of the planning accounts themselves. By classifying costs into "committed" (i.e., fixed, not under management control), "managed" (i.e., discretionary overhead such as management salaries, computer rental) and "product" (i.e., direct and indirect materials and labor), greater precision in attributing variance to particular gross causes is obtained. But there is no evidence of formal diagnosis being employed in this approach. It represents, in a sense, symptomatic planning process control.

Another area in which planning process control has evidently been pursued has been in the military. The classic doctrine of von Clausewitz, which influenced military planning all over the world, required the commander to state formally his goals and concepts and "estimate of the situation" at the outset of the planning process. This forces such explicit statement of his premises and conclusions, that after-the-fact assessment of blame among assumptions, plan, and execution is relatively easy. Confounding the causes of poor performance is avoided.

Planning process control has often been employed in military training; it is not so clear that it occurs under the pressure of actual operations. But the critiques of maneuvers and large scale training exercises frequently focus on the planning process itself as distinct from operations as executed.

Informal, qualitative planning process control is limited in its effectiveness, again because of the discrimination problem. It is one thing to know that an estimate was bad; it is another to know why it was bad. And, because the planning process itself is relatively unstructured, it is difficult to pinpoint the particular subprocesses that were defective.

As we argued earlier, one way to improve the potential of the control process is to move towards more formal planning models. Rigid plan formats (the "five paragraph" military format, and the Westinghouse chart of planning accounts are examples) and specific procedures are steps in this direction. We believe that more detailed and more complex models -- in short,

mathematical or computer models will be of even greater value in forcing explicit assumptions and estimates and organizing the process for controllability.

One example of a trend in this direction is the Apollo Project simulation model recently installed for the NASA Office of Manned Space Flight.²⁸ Because the planning assumptions are, in effect, inputs to a computer program, they are "visible"; because the planning procedure involves explicit recourse to a computer program to examine alternative procedures, it would be possible to obtain a "trace" of the search process. Hence, the raw material, i.e., the basic "measurements" of the planning process are available. But since the planning horizon for APOLLO is long (at least to 1970), the problem is obtaining feedback on planning results. This is a case which calls for partial data analysis.

Another element of our specifications is being implemented at Westinghouse Electric, namely that calling for expanding the corporate data base to include extra-corporate data.²⁹ This also contributes basic measurements for planning process control.

²⁸This system has not yet been publicized. It was designed by Peat, Marwick, Livingston and Company under contract number NASw-1223 and installed this year.

²⁹Described in Burck [7, p. 113].

Some Conclusions on the State of the Art

We conclude from our brief review that there exists no publicized comprehensive realization of intelligent management information systems. We also perceive evidence that intelligence is sought in numerous cases.

Diagnostic control of operating processes seems imminent in restricted environments such as continuous process control and the basic pattern is being established even in such messy discrete process control areas as job shop production control. Employment of well-established statistical methodology holds promise for inductive control at higher operating levels.

Planning process control at present is in some cases performed, but always performed informally. Diagnosis appears to be ad hoc and somewhat political in flavor. It is at best qualitatively based and this, coupled with the generally informal nature of the control process, lead us to suspect that its effects are impermanent even when it is effective. But the general trend towards more formal planning models offer opportunity for greater sophistication in control.

In total, the bits and pieces from which higher level intelligence in management information systems can be fabricated exist. The problem is to assemble these within one organization.

IV. STEPS TOWARDS REALIZATION OF INTELLIGENT INFORMATION SYSTEMS

Introduction

The total realization of what we have advocated requires considerable research and development before it can be accomplished. For example, there exists almost no general theory of diagnosis in particular, and inductive inference in general. On the other hand, we feel that some major improvements in the state of the art could be effected simply by recognizing the value of what we have called intelligence and reorienting the information system to the end of acquiring it. Also, new information technology (including modeling approaches subsumed under operations research such as digital simulation and heuristic programming, as well as "third generation" computer technology enabling real-time data processing and time-sharing) now affords some major capabilities which can be exploited for this purpose.

Therefore (with perhaps more alliteration than accuracy) we have defined our steps toward realization as recognition and reorientation, real-time processing, and research.

Recognition and Reorientation

Given the desire to increase the intelligence of an organization, there are some fundamental steps that can be taken. In our view, the major discrepancy between typical operating process control systems and those which we want is in the explanatory power of the underlying process models. That is to say, conventional operating process control fails to get at the

underlying causes of process variance. Relatively unsophisticated statistical analysis such as analysis of variance and covariance or multiple regression can shed considerable light in this area in many situations. And we suspect that simple classification-association would uncover some of the grosser causes. For example, we recently observed a case in which a salesman's pricing misbehavior was uncovered simply by comparing his customer claims experience with that of the rest of the sales force. The basic statistics showed that his claims were far more frequent and, on the average, larger than those of his colleagues. Deeper investigation uncovered the fact that many of his claims were unrelated to damaged, missing, or sub-standard goods, but were simply his mechanism for granting price concessions to customers (his commissions were not adjusted for claims). We could cite numerous other examples of surprise resulting from attempts to rationalize the causes underlying other performance measurements. It is important that we stress, however, that such rationalization of causes must become part of the information system if intelligent system behavior is to be obtained.

A managerial accounting system has the capability to store raw data and classify them. Classification is done purely on the basis of human intervention, because the system does not have self-organizing characteristics. But given the classification, through a matching process the system can extract differences, which are then reported to management.

A difference by itself, of course, does not mean very much. Although "red" variances (debit balances in manufacturing accounts, for example) are automatically considered undesirable and "black" variances (credit balances) desirable, in reality, much more analysis is necessary beyond this stage. At best these differences may point to a potential problem. The questions that come to mind on observing these accounting variances are mainly of two types: (a) how significant (in a probabilistic sense) are they, and (b) what do they mean?

To enlarge the capacity of the data base and the capabilities of the information systems, one may store in the data base cues for automatic response at the operating level. This response may be purely of the thermo-static control type, as explained before, or the result of elementary analysis performed by the system itself. Such a system does not allow for any ambiguity, in short it is deterministic and inflexible. To generate intelligent behavior, our data base must be capable of resolving ambiguity, and possess understanding and learning capabilities.

Probabilistic distributions can be introduced into the data base and decision rules for determining the probabilistic significance of the observed deviations. So the system can be instructed to sift through the differences, take remedial action on the basis of prestored cues, or else report the significant variations to the manager -- "management by

exception". But we need not stop here. We can also in a Bayesian sense review the models which govern the expectation of system behavior, and also update the relevant probabilistic distributions.

A method of introducing the necessary capability of cause determination in the data base is to store predetermined functional (cause-effect) relationships and explicit decision rules to facilitate their use. Such an arrangement, however, is not very different from thermostatic control and limited in its intelligence. No understanding or inference takes place. We could alternatively instruct the data base in the methods of arriving at hypotheses of cause and effect relationships by itself. This is a more promising avenue because it permits adaptability.

Simple and naive techniques such as statistical variance and covariance analysis, if performed on the accounting variances, can yield useful cause and effect relationships to be stored in the data base for further analysis and testing hypotheses. This type of a system was elsewhere called a functional accounting system [46] and many of its characteristics and prerequisites for implementation have already been discussed [41, 45, 47]. We believe that with the present state of technology and knowledge, the functional accounting system is realizable now. Furthermore, under such a system many facets of the design of organization structures are brought within the purview of the system and resolved analytically for the first time [42].

The major discrepancy in planning process control derives from the informality both of the planning process itself and whatever planning review procedures exist. The initial step, we believe, is to impose some formal requirements on the planning process for purposes of establishing the basic measurements for planning process control. More specifically, we advocate formal planning models, again as part of the information system. It must be granted that formal planning does not imply formal planning process control, but it is the point of departure for setting up the necessary data base for evaluation of performance. Unless there is a systematic method of separating the assumptions and forecasts (the model and its parameters), from the logical deductions as to course of action to be taken therefrom, and then the execution, there is little hope of improvement of the process. Emery [14] has waxed fervent and at length on this subject. We agree with him.

Real-Time Systems: The New Information Technology³⁰

Our contention is that the new computers offer capabilities that enable much easier construction of intelligent management information systems. Both the quality of the data base and the power of the procedures that can be brought to bear on it can be materially improved.

Consider first the now generally available facility for "on-line, real-time" data processing in general and real-time operational control in particular. Since real-time processing implies up-to-the-minute recording

³⁰This section of the paper is a partial synopsis of [8].

of system-wide individual transactions (status changes), it provides uniquely a current, global data base (for operations). In other words, the current status of all operating system activities is known. Furthermore, since all activities are "on-line" to the central processor, access to large-scale computational power can be granted in order to respond to transactions as they arise. This has distinct implications for the situation in which the desired response is in the form of control directions which, recall, are made on the basis of a process model.

Time-sharing is a product of the same technology, being essentially on-line, real-time computation for multiple users. It provides for man-machine interaction in problem solving without creating idle time. The close coupling thus afforded means that there exists a flexible division of labor between man and machine, the man bringing to the process those attributes in which he excels in close cooperation with the superior computational powers of the machine.

Since real-time processing and time-sharing are based on the same technology, they are mutually compatible and both are compatible with conventional "batch" processing it has been noted [8]. Consequently, we can hypothesize the near-term existence of generalized computers which possess real-time processing and time-sharing capabilities in which:

"...whatever permutation or combination of human and machine problem solving attributes is needed can be supplied with data inputs of whatever quality of currency or scope is desired." [8, p. 10]

That these things are good in general is undoubtedly true, but they are particularly useful for the pursuit of intelligence, we assert. In operating process control, we have noted the need for detailed data and functional-temporal association thereof. The global scope and currency of the data in these generalized systems is basic to this requirement. Furthermore, the availability of these data, coupled with the computational power in real-time, means that quite complex and hence potentially more valid process models can be utilized in control. And finally, we have noted that inductive inference is a faculty limited, for the moment at least, largely to human beings, yet it is the fundamental process of increased understanding. Through the new technology, a human can be closely coupled to operating process information; he can monitor the process and exercise his superior capabilities for pattern recognition, abstraction, hypothesis formulation and test -- in short, his inductive powers. He need not await accumulation of evidence; he is on-line to the operating process even though it may be geographically dispersed. As we will observe shortly, this testing capability is very important. Once recognition, reorientation and understanding of the process is established, it may then be introduced into the system itself for increased sophistication and intelligence.

Some of these points can be illustrated by consideration of the application of real-time processing to the job shop production control problems cited earlier. In general, the simulation-based schedules were noted as providing superior control because they provide a more valid model

of the process. Obviously if the model is valid, departures from schedule mean something. But this is true for only a short period after the new schedules are computed. In time, minor deviations from schedule accumulate, machines break down, workers are absent, and as a result, the model (i.e., the schedule) and the real world begin to diverge. This is the "decay" to which we referred previously. In a real-time production control system, detailed decisions on product movement, relating to sequencing and routing of jobs, could be performed by the computer (using the same type of decision rules employed in the simulation). But because the status of the system is continuously updated, the decisions are made on the basis of true current status as opposed to the predicted status used in the simulation approach. Consequently, "decay" would not occur and the resulting control system theoretically would be more effective.³¹ Also, we would have a more useful system for diagnosis because deviations from expected behavior would mean precisely that something other than the model is wrong. Analysis of causes could therefore be undertaken without risk of a wild goose chase, and the fact that investigation can take place immediately provides unparalleled opportunity for accurate reconstruction of the "crime", it is conjectured. But good diagnosis does not stop with crime reconstructions. Its greatest

³¹The relative effectiveness of real-time versus periodic scheduling was tested by Kogan [24]. In the cases studied, the theory was found to be valid. Of course, we must admit that the comparisons are influenced extensively by the expertise of the one who simulates. Even so, since we are dealing with non-deterministic systems, real-time control will outperform controls based on fixed theoretical models.

value lies in its educational aspects. The more confident one is of cause and effect relationships the stronger pattern association and the faster the remedial action.

Diagnosis is, of course, often a more subtle process than is implied by the running down of variations from plan. It often requires "browsing" through historical data, classifying, normalizing, rearranging and the like. The flexible interaction feature of time-sharing provides great convenience and power for so doing. Being able to think between interactions with the computer is at the heart of the concept of "man-computer symbiosis" advanced by Licklider [26] among others.

The advantages of this new technology are perhaps even more marked in the planning process control domain. First of all, planning itself is a natural man-machine process, it has been frequently noted [9, 14]. Simply being able to intertwine the heuristically well-endowed, intuitive and subjective planner with his model offers enormous advantages. But the greater advantage may come in the metamodeling process, that is, modeling the planner's behavior for purposes of ultimate improvement (assuming his use of a computer model). Capturing the planner's "trace", the detailed sequence of steps he takes in arriving at a decision, is quite possible. Given the planner's cooperation, it is simply a question of obtaining the hard copy transcript of his session with the computer model, for example. And, of course, the diagnostician is equipped with a unique linkage to the process he is attempting to understand.

Exploitation of the power of these generalized systems for operating process control, planning, and planning process control has been the subject of research by Carroll and colleagues at Project MAC [10].

Research in Intelligent Information Systems

When we view where we stand in relationship to our goal of intelligence, we realize just how little is known about the techniques, the economics, and, broadly, the phenomenon of intelligence.

We list below a few areas of research which we feel represent promising starting points for improvement in this regard.

Data Base and Information Systems for Intelligence

We need better understanding of the data required for improved understanding. We have stated the general specifications for detailed and associative data; but there are numerous questions begged by this, such as what detail and what means for association. In short, we need some operational specifications (subjected to economic analysis hopefully) and some demonstrably better mousetraps in the data structure domain. There is a dilemma involved in the question of what detail, for example. One simply cannot specify a priori what detail or even what variables to measure until unexplained problems (or successes) occur, and hypotheses are generated. What needs to be established is the usefulness, for diagnostic purposes, to provide guidelines on collection of possibly relevant data (as opposed to already known relevant data). In short, some theory is needed.

Some general research in the structure of associative data bases and the procedures for exploiting this association is in process by Zannetos and Sahin [44]. This can be described (speculatively) as follows.

The data base requirements for an associative information system will be mainly the same as those of a functional accounting system, with one major difference. Instead of using raw data as the indecomposable modules for storage, manipulation, and causal association, we will now use patterns, or configurations of data. The faculties of understanding, we might even say "consciousness", loom into prominence and somehow must be captured and incorporated in the data base.

To get at the question of procedure, assume that we have an organization with well-established objectives and a dominant (i.e., chosen) plan to accomplish them. Given the dominant plan, we assume that the organization will be able to specify the operations that are necessary to achieve its objectives. Now for each dominant plan there must be a given configuration (pattern) of resource utilization, at least on an a priori basis. With the resource configuration established, a dominance ranking of these resources can be made in terms of a one-dimensional index. Such ranking may be in terms of opportunity costs or loss functions. We are only interested in the dominant plans and resources, and in proximate ordinal rankings. (The hypotheses which the system will generate and the search which will follow for testing will compensate for such approximations.) Furthermore, we are only interested in probabilistic associations.

The next requirement is the association of resources, at the point of acquisition, with the various (major) attributes of such resources. These sets of attributes are given a temporal index and also contain entries indicating the major physical characteristics of resources (among which are cost and capacity information). The attributes of resources are ranked once again according to dominance which obviously is dictated by the dominant plan.

Each one of the resources in the dominant configuration, no matter what its ranking therein, may be the dominant resource in another configuration of subordinate resources, as well as non-dominant member of other configurations. We thus have a hierarchy of associations both vertical and horizontal.

With the above as a brief description of the system, let us now look at hypothesis generation at the operating level, because this is one of the greatest attributes we wish to impart to the information system. The signals which trigger hypothesis generation, are of at least three kinds. They may originate in:

1. The difference between resource utilization (both quantities and attributes used) as specified in the prior dominant plan and as reflected in operations.
2. The dominant resource configuration of a proposed plan, if it does not use the most dominant characteristic of each of the resources included in the plan. (If the new plan, after search, is still found to be dominant then an updating of the resource-attribute vectors will be necessary.)

3. The presence of "slack" in some dominant resource which will necessitate a change in the opportunity cost of this resource and a temporary change in the dominance rank. (The system scans for slack in capacity starting from the most dominant resource downward.)

Once a signal is received, on the basis of its content, the system immediately associates at least two patterns with it: the highest hierarchical pattern where the resource appears, whether in a dominant role or not, and the pattern in which it is the most dominant resource. Now the search begins for term by term comparisons of the prior patterns (plan) and those derived from operations or the proposed plans, and hypotheses are tested. Depending on the results of these tests the descriptive sets of resource attributes may be rearranged and assigned new temporal designation. Also, the data base is hierarchically reorganized.

In addition to its diagnostic properties, this system also holds promise in providing information for prognostic purposes. By studying the intertemporal changes in the resource-attribute vectors, the system may generate hypotheses and test for the existence of patterns of relationships which can be used for planning process control. This is a task for which the system needs to interact with man.

For a system such as this to operate efficiently, we cannot obviously depend on "brute force" or exhaustive sequential search, because we then get into immense combinatorial problems. We suspect that we must use, therefore, parallel search techniques or some hybrid system.³² As for the cues that trigger pattern retrieval and association, they must not refer to locations of stored messages but to the content. Finally, we noticed that there is a need for some hierarchical organization of the data base with distributed logic. This relative decentralization allows flexibility for learning and self-organization, but also necessitates functional association of the various modules of the data base. The problem of deciding how much logic is to be distributed and where is not an easy one to solve. We believe, nonetheless, that it is not unlike other organizational problems, so the theory and techniques suggested elsewhere for aiding in the design and evaluation of the organization structure are also applicable in this case [40, 41, 43].

Theories of Diagnosis and Decision Making

We have noted that the "metamodeling" problem of planning process control necessitates modeling the planning process itself. But planning is a decision process, so this amounts to modeling a decision maker. This has been an active area of research for some years now, notably by students of Simon and Newell such as Clarkson [12]. Other approaches have been studied

³²Selfridge and Neisser [37] have commented on the relative merits of the two search strategies.

by Bowman [5] and his students. However, no direct research has been directed to modeling the type of modeling process involved in planning and we think this would be useful.³³

Moreover, diagnosis, in the sense that we have employed the term, is a poorly understood process at best. There is much work going on in medical diagnosis, but unfortunately for our purpose this is what might be called "discriminatory diagnosis" in which the relationship of symptoms to diseases is taken as known (probabilistically) rather "inductive diagnosis" in which no such relationship is available.³⁴ However, work in medical diagnosis will undoubtedly provide some general insights. Particularly promising is the research of Gorry [20] who is attempting to create a general, diagnostic model ("general" means environment independent -- applicable to sick people, sick cars, sick computer programs). His emphasis is on discriminatory diagnosis, but we suspect that the data structures and much of the logic of his procedure are applicable to the less well-structured inductive diagnostic problem as well.

In addition to understanding individual decision behavior, we have noted that planning is often a group process; it is performed by a "team". Team decision making is not well-understood. The pioneering work of Radner [32]

³³Newell and Simon [29] did incorporate a "planning" mechanism in their "General Problem Solver", it should be noted, but their type of planning and ours are only generically related.

³⁴What must be supplied in inductive diagnosis is the hypothesis of relationship, a task we have already allocated to the man in the man-machine partnership.

Marshak [27] and Kriebel [25] brings some organization to this area and the recent work of Clarkson [11] in descriptive (computer simulation) approaches to group decision making is directly relevant to the problem. The coming general availability of time-sharing, which enables group cooperative, interactive problem solving and monitoring, should greatly facilitate research in this area.

V. SUMMARY AND CONCLUSIONS

We have specified some features of information systems which are capable of increasing management's understanding of its environment and its rationality in coping with it -- its intelligence. In so doing, we have focused our attention first on operating process control, in which intelligence is increased by causal induction, diagnosis of the factors which underlie process behavior. This is conveniently framed as improving the model of the process. We then discussed the higher level problem of planning process control which we noted was typical of higher order intelligence problems. This too involves model improvement, but in this case the model is of what is a modeling process itself, in part.

We have reduced our discussion to size by ignoring some aspects of system design. For example, we have ignored the general dimension of information availability. There exist in this area several issues, to wit: Should information relating to detailed performance of lower level organizational subunits be freely accessible to higher level managers? Should parallel organizational units share data on their status and performance? These are real issues which are related in part to "managerial style" but they also impinge on organizational intelligence. The "multiple split personality" aspect of organizational behavior is, we suspect, intimately linked to the question of dissemination of information. We have assiduously avoided the

question of who should perform diagnosis -- the controlled entity, the controller, or even some disinterested party. This question we leave to the social psychologists.

Another area that we have ignored encompasses the perennial issues of "cost and value" of the generated information. No doubt, trade offs must be established between the cost of the system, detail and purity of information, reality of representation among other features, and the objective as well as the often subjective utility of the results. All these issues we chose to leave outside the purview of this presentation for reasons of expediency and without prejudice.

Within the scope of the general problem we have attacked, we conclude that increased organizational intelligence is possible and greatly to be facilitated by new advances in information technology. Perhaps the greatest progress can be made, however, by simply recognizing that increasing intelligence is a legitimate goal for information systems design, and that there are some straightforward steps which can be taken towards that goal.

References

1. Anthony, Robert N. Planning and Control Systems: A Framework for Analysis. Boston: Division of Research, Graduate School of Business Administration, Harvard University, 1965.
2. Armer, Paul. "Attitudes Towards Intelligent Machines," in [17], 389-405.
3. Bellman, Richard. Adaptive Control Processes. Princeton: Princeton University Press, 1961.
4. Black, William A. "Cost Models for Management Control," unpublished masters' thesis, Sloan School of Management, Massachusetts Institute of Technology.
5. Bowman, E. H. "Consistency and Optimality in Managerial Decision Making," in J. L. Muth and G. L. Thompson, eds., Industrial Scheduling. Englewood Cliffs: Prentice-Hall, 1963, 99-112.
6. Brown, Robert G. Smoothing, Forecasting, and Prediction of Discrete Time Series, Englewood Cliffs, Prentice Hall, 1963.
7. Burck, Gilbert, et. al. The Computer Age and Its Potential for Management. New York: Harper and Row, 1965.
8. Carroll, Donald C. "Implications of On-Line, Real-Time Systems for Managerial Decision Making," Paper presented at the Research Conference on the Impact of New Developments in Data Processing on Management Organization and Managerial Work, Sloan School of Management, Massachusetts Institute of Technology, March 29-30, 1966. (Also Working Paper 165-66, Sloan School of Management, MIT).
9. Carroll, Donald C. "Man-Machine Cooperation on Planning and Control Problems," Paper presented at the International Symposium on Long Range Planning for Management, UNESCO, Paris, September, 1965. (Also Working Paper 145-65, Sloan School of Management, MIT).
10. Carroll, Donald C. "Simulation Research in On-Line, Real-Time Systems," Paper presented at the TIMS Eastern Meetings, Rochester, N. Y., October 15, 1965. (Also Working Paper 164-66, Sloan School of Management, MIT).

11. Clarkson, Geoffrey P. E. "Decision Making in Small Groups: A Simulation Study," Working Paper 194-66, Sloan School of Management, Massachusetts Institute of Technology, 1966.
12. Clarkson, Geoffrey P. E. Portfolio Selection: A Simulation of Trust Investment. Englewood Cliffs: Prentice Hall, 1962.
13. Emery, James C. "An Approach to Job Shop Scheduling Using a Large-Scale Computer," Industrial Management Review, III (Fall 1961), 78-96.
14. Emery, James C. "The Planning Process and Its Formalization in Computer Models," in Joseph Spiegel and Donald Walker, eds., Information System Sciences: Proceedings of the Second Congress. Washington: Spartan Books, 1965, 369-390.
15. Evans, Marshall K. "Profit Planning," Harvard Business Review, 37 (July-August, 1959) 45-54.
16. Ezekial, Mordecai and Karl A. Fox. Methods of Correlation and Regression Analysis. Third Edition. New York: Wiley, 1959.
17. Feigenbaum, Edward A. and Julian Feldman, eds. Computers and Thought. New York: McGraw-Hill, 1963.
18. Forrester, Jay W. "A New Corporate Design," Industrial Management Review, VII (Fall, 1965), 5-14.
19. Forrester, Jay W. Industrial Dynamics. Cambridge: MIT Press, 1962.
20. Gorry, Anthony. "The Diagnostic Problem and Process," Ph.D. Dissertation in process, Massachusetts Institute of Technology.
21. Graybill, Franklin A. An Introduction to Linear Statistical Models. New York: McGraw-Hill, 1961.
22. Harvey, Edward H. "A Study of the Formulation of Operating Cost Models for Engineering of Economic Analysis in the Telephone Industry," unpublished masters' thesis, Sloan School of Management, Massachusetts Institute of Technology, 1966.
23. Holland, James G. and B. F. Skinner. The Analysis of Behavior. New York: McGraw-Hill, 1961.
24. Kogan, John N. "Real-Time Sequencing versus Periodic Scheduling in a Job Shop," unpublished masters' thesis, Sloan School of Management, Massachusetts Institute of Technology, 1966.

25. Kriebel, Charles H. "On Normative Mathematical Models of Information Systems," Working Paper 38-63, Sloan School of Management, Massachusetts Institute of Technology, 1963.
26. Licklider, J. C. R. "Man-Computer Partnership," International Science and Technology, (May, 1965), 19 et seq.
27. Marschak, Jacob. "Elements of a Theory of Teams," Management Science, 1 (January, 1955), 127-137.
28. Minsky, Marvin. "Steps Toward Artificial Intelligence," in [17], 406-452.
29. Newell, Allen and H. A. Simon. "GPS, A Program that Simulates Human Thought," in [17], 279-293.
30. Newell, Allen and H. A. Simon. "The Simulation of Human Thought," Current Trends in Psychological Theory. Pittsburgh: University of Pittsburgh Press, 1961, 152-179.
31. Pounds, William F. "The Process of Problem Finding," Working Paper 147-65, Sloan School of Management, Massachusetts Institute of Technology, 1965.
32. Radner, R. "The Application of Linear Programming to Team Decision Problems," Management Science, VI (January 1959), 143-150.
33. Raiffa, Howard and Robert Schlaifer. Applied Statistical Decision Theory. Boston: Division of Research, Graduate School of Business Administration, Harvard University, 1961.
34. Roethlisberger, F. J. and W. J. Dickson. Management and the Worker. Cambridge: Harvard University Press, 1939.
35. Schlaifer, Robert. Probability and Statistics for Business Decisions. New York: McGraw-Hill, 1959.
36. Selfridge, Oliver G. and Ulric Neisser. "Pattern Recognition by Machine in [17], 237-250.
37. Steinhoff, Harry W. Jr. "Daily System for Sequencing Orders in a Large Scale Job Shop" in Elwood S. Buffa, ed., Readings in Production and Operations Management. New York: Wiley, 1966, 155-165.
38. Trilling, D. R. "The Use of a Job Shop Simulator in the Generation of Production Schedules," AFIPS Conference Proceedings Vol. 26: 1964 Fall Joint Computer Conference. Baltimore: Spartan Books, Inc., 1964, 277-290.

39. Turing, A. M. "Computing Machinery and Intelligence," in [17], 11-35.
40. Zannetos, Zenon S. and Donald C. Carroll, "Information-Technological Pressures in Organizational Design," Working Paper 219-66, Sloan School of Management, Massachusetts Institute of Technology, 1966.
41. Zannetos, Zenon S. "Managerial Information Systems for Planning and Control," Working Paper 210-66, Sloan School of Management, Massachusetts Institute of Technology, 1966.
42. Zannetos, Zenon S. "Measuring the Efficiency of Organization Structures: Some Implications for the Control System of the Firm," Working Paper 117-65, Sloan School of Management, Massachusetts Institute of Technology, 1965.
43. Zannetos, Zenon S. "On the Theory of Divisional Structures: Some Aspects of Centralization and Decentralization of Control and Decision Making," Management Science. 12 (December 1965), 49-69.
44. Zannetos, Zenon S. and Kenan Sahin. "Some Thoughts on the Design and Implementation of Associative Information Systems for Planning and Control," Working Paper 220-66, Sloan School of Management, Massachusetts Institute of Technology, 1966.
45. Zannetos, Zenon S. "Standard Costs as a First Step to Probabilistic Control: A Theoretical Justification, an Extension and Implications," The Accounting Review, XXXIX (April 1964).
46. Zannetos, Zenon S. "Toward a Functional Accounting System: Accounting Variances and Statistical Variance Analysis," Industrial Management Review, 7 (Spring 1966), 71-83.
47. Zannetos, Zenon S. "Toward Intelligent Management Information Systems," Working Paper 155-65, Sloan School of Management, Massachusetts Institute of Technology, 1965.
48. Zschau, Edwin V. W. "A Technique for Comparing the Performance Results of Operating Units Within a Decentralized Organization." Working Paper No. 41, Graduate School of Business, Stanford University, 1964.

21

Date Due

DEC 11 '75	FEB 1 1976	
DEC 22 '75		
JAN 20 '76		
FEB 18 '76		
DEC 27 '75		
SEP 15 '76		
MAY 21 '76		
SEP 10 '76		
MAY 25 '79		
MAY 21 '82		
		Lib-26-67



175:1.1
HD28 M.I.T. Alfred P. Sloan
.M414 School of Management
Nos.211-66 Working Papers.
Nos.224-66 Nos.211-66 to Nos.224-66



212-66



213-66

PASSEMENT



214-66

SIGN this card and present it with book
at the CIRCULATION DESK.



215-66

MASSACHUSETTS INSTITUTE OF
TECHNOLOGY LIBRARIES



216-66



217-66

218-66



4-66



1-6



